arXiv:1703.04489v1 [cs.CL] 13 Mar 2017

Reinforcement Learning for Transition-Based Mention Detection

Georgiana Dinu, Wael Hamza and Radu Florian
IBM T.J. Watson Research Center
1101 Kitchawan Rd, Yorktown Heights, NY 10598
whamza|gdinu|raduf@us.ibm.com

Abstract

This paper describes an application of reinforcement learning to the mention detection task. We define a novel action-based formulation for the mention detection task, in which a model can flexibly revise past labeling decisions by grouping together tokens and assigning partial mention labels. We devise a method to create mention-level episodes and we train a model by rewarding correctly labeled complete mentions, irrespective of the inner structure created. The model yields results which are on par with a competitive supervised counterpart while being more flexible in terms of achieving targeted behavior through reward modeling and generating internal mention structure, especially on longer mentions.

1 Introduction

Mention detection is the natural language processing task of identifying and assigning labels to occurrences of entities in text, according to a predefined type-system. For example the following sentence (from Pradhan et al. [2013]): *Monday night, Colin Powel, Robert Downey Junior and Sharon Stone.* contains three mentions of type PERSON: Colin Powel, Robert Downey Junior, Sharon Stone and a mention labeled TIME: Monday night. As this is a joint chunking and labeling task, the standard approach is to formulate it as a sequence labeling problem, in which the label of a mention word encodes both the type and the position within the mention. In the previous example, a correct PER classification within the BIO encoding scheme is $\text{Robert}_{\text{B}-\text{PER}}$ $\text{Downey}_{\text{I}-\text{PER}}$ $\text{Junior}_{\text{I}-\text{PER}}$.

While performance for this task is relatively high for short mentions, longer mentions are considerably more difficult to label correctly. Consider for example the sentence: *The George Washington Bridge is owned and operated by the New York and New Jersey Port Authority.* A standard annotation for this sentence labels [George Washington Bridge]$_{\text{FACILITY}}$ and [New York and New Jersey Port Authority]$_{\text{ORG}}$ as mentions. Such examples pose difficulties to models due to the inherent bias created by, for example, New York mostly occurring as GPE. This is still a problem within neural network architectures (Andor et al. [2016]). In order to mitigate the label bias problems (Bottou [1991], Lafferty et al. [2001b]), models may require significant look-ahead and increased capacity (to encode that e.g. New is an ORG when followed by Port Authority or that George in the context of Bridge is a FACILITY) or global normalization with e.g. CRF objectives.

Alternatively, having the capacity to generate internal structure and update labels as this structure is built may allow a model to more naturally revise labeling decisions. For example, a correct label can be obtained by initially labeling George Washington as a PERSON and then composing [PERSON Bridge] into FACILITY, both decisions corresponding to frequent enough patterns which a model should be able to learn more easily from training data. For this purpose, we propose a transition-based formulation which can incrementally construct mentions as labeled binary trees and which is trained in a reinforcement learning setting. Specifically, the labeling is obtained as the output of a sequence of actions, allowing it to either label an incoming token or compose the 2 nodes on the top of the current stack and assign it a (potentially different) label. As inner structure is not usually

annotated, and hence cannot be learned directly, a reinforcement learning setting is the natural choice for learning and predicting it. In this paper, we propose a Q-function approximating network which positively rewards a sequence of actions which predicts a correct mention tree root. Crucially, we do not assign any reward to partial decisions in order to encourage the model to freely learn to label incomplete mention spans.

The remainder of the paper is organized as follows: Section 2 discusses relevant literature. Section 3 introduces our proposal for mention detection through reinforcement learning while Sections 4 and 5 present the experimental evaluation and conclude the paper.

2 Related work

Mention detection and named entity recognition[1], have been extensively studied in the literature, due to their usefulness to end-user tasks (such as question answering, converting unstructured information into structured data, etc). Datasets for named entity recognition have been created since the mid 90s, with the MUC-6 dataset, then continued with the CoNLL datasets (Tjong Kim Sang [2002], Tjong Kim Sang and De Meulder [2003]), the ACE dataset (NIST [2003]), and the OntoNotes dataset (Pradhan et al. [2013]). The problem has also been studied across languages (English, Arabic, Chinese, Spanish, German, Dutch, and Hindi are among the most studied languages). Most of the approaches treat the task as a sequence labeling task, where each word gets assigned a tag corresponding to whether it is a non-mention word (usually, "O"), starts a mention or continues a mention. This corresponds to a reversible mapping from the original chunks to word-based labels (this BIO scheme was first described by Ramshaw and Marcus [1994] for noun phrase chunking), and leads to standard sequence-based classification models such as maximum entropy Borthwick et al. [1998], Florian et al. [2004], CRFs ([Lafferty et al., 2001a]) and neural-network based approaches ([Collobert et al., 2011]). Approaches that attempt to produce the chunks directly have also been investigated, such as the one by Daumé and Marcu [2005].

A transition-based model for NER as an alternative to the standard encoding has been introduced in Lample et al. [2016], where it is shown to obtained state-of-the-art results among comparable models and it is only outperformed by a globally normalized LSTM model. The formulation is inspired by dependency parsing in which the labeling is the result of a sequence of *transitions* which at each step decide whether to mark an incoming word as non-mention, start/continue a mention or end a mention and assign a label. Under this formulation, for each labeled sequence there is an unique sequence of transitions which generates the correct annotation and therefore it is natural to remain within a supervised learning setting.

Reinforcement learning for sequence labeling problems in NLP (including Named Entity Recognition) has been proposed in Maes et al. [2009]. Unlike our proposal, the action space closely follows the supervised NER formulation consisting of token-level labels using the BIO encoding, assigned either in a left-to-right order or in a free-order fashion. Rewards are given for each episode or for each action (positive if the label assigned is the same as the gold standard). They obtain competitive results when using left-to-right order and action-level rewards which outperform the more flexible free-order and episode reward alternatives.

Our work comes closer to that of Jiang et al. [2012] in which the authors propose a hybrid reinforcement/apprenticeship learning algorithm for transition-based dependency parsing in order to obtain a better speed-accuracy trade-off. As a policy gradient algorithm as well as an apprenticeship learning approach fail to reach competitive performance, the authors propose an oracle-infused hybrid which during exploration uses the oracle to generate the correct action with fixed probability.

3 Mention Detection and Reinforcement Learning

We formulate the problem as a Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, \mathcal{R})$, where $\mathcal{S}$ is the state space with discrete actions $\mathcal{A}$ and a reward function $\mathcal{R} : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$. We learn the state-action value function Q and our policy π is a greedy policy in Q. The following section describes the formulation of the transition-based mention detection problem as an MDP.

[1]The main difference between the two is that mention detection has a larger scope: besides named mentions, it also identifies nominal and pronominal mentions.

3.1 MDPs for mention detection

Transition Based Mention Detection The mention detection agent scans the given sentence from beginning to end making labeling decisions, given by the action space, sequentially and one word at a time. During learning, the agent is given rewards based on the gold annotation in the training data. The decision process terminates when the process consumes the entire sentence.

States We represent the current state using a stack and a buffer (using the terminology in Lample et al. [2016]). The stack contains the sequence of tokens which have already been labeled, together with their labels, while the buffer consists of the words left to be processed. Table 1 contains an example of a buffer and stack values for a transition sequence example. The neural network approximating the Q function uses features extracted from the current state, typically some features of the current word (top of the buffer) and several past tokens (the top elements of the stack). The specific input representation used is detailed in Section 4.1.

Action space Given a label set $\mathcal{L}$ we define an action space of size 2x$|\mathcal{L}|$:

$$\mathcal{A} = \{\text{Shift}_{\text{LABEL}}, \text{Reduce}_{\text{LABEL}} | \text{LABEL} \in \mathcal{L}\}$$

Shift actions move a token from the top of the buffer to the stack and assign it the class LABEL. Reduce actions group together and label the top 2 elements of the stack, which can be multi-word sequences themselves. There is no restriction on the $\text{Reduce}_{\text{LABEL}}$ to be consistent with the classes already assigned to the stack elements, i.e. the model can decide to re-label elements - for instance, joining a PERSON group with an O group to form a FACILITY group. A sequence of actions is uniquely mapped to a label sequence and it is read off the stack at the end of an episode. Table 1 contains an example. Non-mention (O labels) are treated like any other labels with the exception of Reduce-O actions which are not allowed.

Table 1: Transition sequence example for *George Washington Bridge*

Stack	Buffer	Decision
$\emptyset$	George Washington Bridge	Shift-PER
$\text{George}_{\text{PER}}$	Washington Bridge	Shift-PER
$\text{George}_{\text{PER}}$ $\text{Washington}_{\text{PER}}$	Bridge	Reduce-PER
$[\text{George Washington}]_{\text{PER}}$	Bridge	Shift-O
$[\text{George Washington}]_{\text{PER}}$ Bridge_{O}	$\emptyset$	Reduce-FAC
$[\text{George Washington Bridge}]_{\text{FAC}}$	$\emptyset$	-

Episodes and rewards At test time, the episode is defined on the sentence level, that is, the agent starts at the beginning of the sentence and continues to make decisions until the entire sentence is consumed. However, during learning, we explore two types of episodes: *sentence* episodes (used at test time) and *mention* episodes. For sentence episodes the reward can simply be the F-score of the sentence (or a function thereof). For mention-level episodes, the agent starts the episode from the beginning of a mention (not necessarily at the beginning of the sentence), and we assign positive rewards if the entire mention is correctly labeled, negative otherwise. We do not assign any reward for intermediate decisions within the mention boundaries. In one of our experiments, however, we experimented with assigning partial reward in cases where the agent forms correct partial mentions, where partial mentions are considered correct if they occur in the training data a minimum number of times. The reader can find details about this in Section 4. The mention-based episodes terminate either when the agent labels the mention correctly or when it makes an error that cannot be corrected. We call these states terminal. Note that words that are not part of mentions are treated as single-word mentions with "O" labels.

3.2 Learning

As mention earlier, we approximate a state-action value function $Q(S, A)$ using a neural network $Q(S, A; \theta)$ that takes the representation of a state S as input and outputs a value for each action $A \in \mathcal{A}$. We use Temporal Difference (TD) learning to learn θ, the set of parameters used in the

neural network. We specifically experimented with both SARSA and Q-learning TD algorithms (Sutton and Barto [1998]). The parameters θ are randomly initialized and the agent uses an ϵ-greedy policy to explore. During exploration, we update the parameters θ by stochastic gradient descent (SGD) using gradients of the following loss function (Q-learning version):

$$L(\theta) = (r + \gamma \max_{a'} Q(s', a'; \theta^-) - Q(s, a, \theta))^2$$

where r is the reward the agent observes after taking action a from state s, arriving at state s', and γ is the reward discounting factor. The *target* $Q(.; \theta^-)$ is a copy of Q which is updated less frequently to stabilize the learning process as in Schaul et al. [2016]. In order to keep the SGD update samples relatively independent, we compute the gradient of the loss function above for a batch of K agents exploring different sentences (or mentions in the mention-based episodes). For each time step in the learning process, we randomly choose K episodes from a pool of concurrent episodes N. When any of the concurrent episodes terminates, it is restarted by randomly selecting a new starting point from the training data. The algorithm for one-step Q-learning is given in Algorithm 1.

Algorithm 1 Multi-agent one-step Q-learning algorithm for transition-based mention detection

Input: D: training data, K: mini-batch size, N: number of concurrent learning episodes, T: total learning steps, α: learning rate, T_u: θ^- update period,γ: reward discounting factor
Initialize θ for $Q(s, a; \theta)$ arbitrarily
Set $\theta^- \leftarrow \theta$
For i = 1 to N:
    Initialize s_i by randomly choosing a starting point from D
For t = 1 to T:
    B = Choose K episodes from N randomly
    For j = 1 to K:
        Set $i \leftarrow B[j]$
        $s_j \leftarrow s_i$
        Choose action a_j from state s_j according to ϵ-greedy policy derived from Q
        Take action a_j and observe state s'_j and reward r_j
        If s'_j is terminal:
            Set $R_j \leftarrow 0$
            Set s_i to a starting point randomly chosen from D
        Else:
            Compute $R_j \leftarrow \gamma \max_{a'} Q(s'_j, a', \theta^-)$
            Advance episode i by setting $s_i \leftarrow s'_j$
    Compute $\nabla\theta = \frac{1}{K} \sum_{j=1}^{K} \frac{\partial [r_j + R_j - Q(s_j, a_j, \theta)]^2}{\partial \theta}$
    Update parameters $\theta \leftarrow \theta + \alpha \nabla \theta$
    If t mod T_u == 0:
        Update $\theta^- \leftarrow \theta$

4 Experiments

We use two English and one Spanish mention detection data sets that vary in terms of size as well as the type system used (Table 2).

4.1 Models

Episodes and rewards We initially set episodes to consist of individual sentences and we assign positive rewards only at the end of a sentence. Given a state-action pair (s_i, a_i) at time step i, the

Table 2: Data set statistics: number of mention types (excluding O) and number of tokens.

Data set	#Labels	#Train	#Dev	#Test
CoNLL En Tjong Kim Sang and De Meulder [2003]	4	232K	58K	53K
CoNLL Spa Tjong Kim Sang and De Meulder [2003]	4	270K	54k	52k
OntoNotes En Pradhan et al. [2013]	18	1778K	271K	186K

reward is 0 if s_{i+1} is not terminal, i.e. the buffer is not empty, otherwise the reward is the F-score or the accuracy (scaled to the [0:10] range) of the labels generated by the $a_{1:i}$ action sequence.

Delaying the reward until the end of the sentence might strain the model unnecessarily as mention-level assignments can be evaluated independently. However, since the model can flexibly change the labels assigned as well as the extent/label of a past mention, it is not immediately clear how to reward actions before a terminal state is reached. To overcome this shortcoming, we train on mention-level episodes, where the boundaries are obtained from the gold annotation. Specifically, each mention and each non-mention word become individual episodes (of length 1 in the case of non-mentions). An episode is finished either when an unrecoverable error has been made or all the tokens of the mention have been processed. We assign a non-0 reward only when the episode is over, equal to ± 1 scaled by the length of the mention.

Learning We use one-step learning in the case of mention-based episodes and a larger TD step for sentence-based episodes. We update the parameters of the target network (θ^-) every 5000 time steps. We use a budget of 3 million (T=3,000,000) for OntoNotes and 1 million for the much smaller CoNLL data sets. In both cases we use a batch size of 10 (K=10).

Baselines We use as baseline a feed forward NN trained in a supervised fashion optimizing the word-level log-likelihood. We use the same network for all data sets, trained using SGD with a linearly decaying learning rate (0.5:0.001) over 15 epochs.

NN-RL architecture Both the Q-function and the supervised baseline use the same following network architecture. The input is the concatenation of the target and context words (symmetric window of size 4) to which we add vectors for three standard features: dictionaries, capitalization flags and character-level representations (concatenation of the last hidden states of forward and backward LSTMs) as well as the previous 2 assigned labels. The word vectors are initialized with 300-dimensional pre-trained embeddings build on a concatenation of Gigaword, BOLT and Wikipedia (totaling 6 billion tokens). Both the additional feature vectors as well as the word vectors are fine-tuned during training (i.e. error is back-propagated to the input representation). We use one hidden layer of size 2000 and sigmoid as its activation function (relu for Q function model) and use 0.5 dropout on the two final layers.

We build a Q-function network by mapping the features used in this standard NN architecture to the new model. More precisely, features that are functions of the target word operate on the top word of the buffer, while the left context is represented as features of the stack. We use a max-pooling 2-gram convolution layer to obtain a fixed-length representation for the words grouped in one stack element.

Hyper-parameter tuning Additional datasets (not used in this paper) have been used to determine robust settings for the NER/mention detection tasks, determining for example that SGD with decaying rate is better than Adagrad/Adadelta, 1 large hidden layer performs better than deeper architectures and modeling context as word concatenation outperformed the use of LSTMs. When available, hyperparameters reported in previous work as optimal (e.g. dimensionality of character-level LSTMs, dropout value) were employed.

4.2 Results

Table 3 reports F1 scores on the three data sets including 2016 state-of-the-art results, the supervised baseline system (SL) and reinforcement learning system (RL) with mention episodes.

SARSA and Q-learning perform similarly. The best results are obtained with mention episodes and the RL models perform similarly to the supervised counterpart on CoNLL English, and drop up to 1 to 2 F points on OntoNotes and CoNLL Spanish respectively. As expected, sentence episodes (omitted in Table 3) take much longer to converge and lead to overall poorer performance: Using a curriculum based on sentence length the models reach 87 F score on CoNLL English and only 56 on the more difficult Spanish data set, while using no curriculum halves the F scores.

Table 3: Mention detection F1 scores: State of the art, supervised baseline (SL) and Sarsa/Q-learning using mention-level episodes (RL). *-with tuned hyper-parameters.

	CoNLL English		OntoNotes		CoNLL Spanish	
Model	Dev	Test	Dev	Test	Dev	Test
Lample et al. [2016]	-	90.94	-		-	85.75
Chiu and Nichols [2015]	94.03*	91.62	84.57*	86.28	-	-
SL	93.06	89.85	86.39	85.59	83.82	84.64
RL-Sarsa	93.22	89.28	85.15	85.35	81.29	83.61
RL-Qlearn	93.25	89.17	85.23	84.67	81.72	84.35

Sub-mentions We further analyze the structure and the partial labels that the RL model assigns to investigate whether sub-mentions receive labels which are consistent with other occurrences as individual mentions. Specifically, for each mention, we extract all its occurrences, if any, as a subsequence of a larger mention in the development and test data (for example Hong Kong and Hong Kong Disneyland). We further count how many times the original mention is identified as a chunk ([[Hong Kong] Disneyland] vs Hong [Kong Disneyland]]) as well as how many times the assigned label is consistent with the original label (e.g. [[Hong Kong]$_{\mathrm{GPE}}$ Disneyland] vs. [[Hong Kong]$_{\mathrm{ORG}}$ Disneyland]).

We examine the previous RL model as well as an additional variant denoted (RL+partial rewards) which gives partial rewards before mention episodes are complete. More precisely, this model encourages "reasonable" sub-mention structures by rewarding partial decisions which lead to sub-mentions that have been seen in training.

Table 4: Labels assigned to sub-mentions broken-down by mention size. Chk: percentage of times the sub-mention is a chunk of the larger mention (Chk rnd: probability of random). Corr. lbl: Percentage of times the correct label is assigned to it. Chance: 0.08

		RL		RL+partial reward		
Sz	Chk rnd.	Chk	Corr. lbl	Chk	Corr. lbl	Correct example
1	1.0	1.0	0.11	1.0	0.65	[two [hours]$_{\mathrm{TIME}}$]$_{\mathrm{TIME}}$
2	0.44	0.46	0.51	0.65	0.88	[[New York]$_{\mathrm{GPE}}$ Stock Market]$_{\mathrm{ORG}}$
3	0.33	0.41	0.71	0.60	0.93	[the [Sino - Japanese]$_{\mathrm{NORP}}$ war]$_{\mathrm{EVENT}}$
4	0.31	0.35	0.68	0.62	0.92	[end of [August of last year]$_{\mathrm{DATE}}$]$_{\mathrm{DATE}}$

Table 4 shows the results broken down by (sub-)mention size. The mention words are grouped together only marginally above chance level. When words are grouped correctly, the labels assigned are correct more than half of the time. This accuracy is much lower for mentions of size 1 (11%) where we observe the assigned label to be meaningless most of the times, e.g. [[China]$_{\mathrm{TIME}}$ Daily]$_{\mathrm{ORG}}$. We hypothesis this is happening because the model is using its look-ahead capabilities to assign correct labels and sees no benefit in making reasonable partial label assignments. The additional partial rewards used with the RL+partial model lead, as expected, to much more coherent sub-mention structures, both in terms of correct chunking and of labels assigned. Surprisingly though, this does not improve overall results, with RL+partial scoring $\approx$ 1F point bellow RL on the OntoNotes development dataset.

Label bias errors We devise a similar method to identify errors which might be caused by label bias. More precisely, we extract cases such as New York Stock Market incorrectly tagged (correct label is ORG), possibly due to the [New York]'s strong GPE bias. The larger mention might be incorrectly identified in a number of ways. We hypothesize that label bias may play a role if either the

entire mention or any prefix of it is tagged as GPE (Geographical and Political Entity), resulting in the extent being misclassified, such as in: [New York]$_{\text{GPE}}$ [Stock Market]$_{\text{ORG}}$. We are again using the development and test portions of OntoNotes and investigate the system output of the supervised and of the Q-learning transition-based models. The results are in Table 5.

Table 5: Number of potential label bias errors made by each model, broken down by size of inner mention (0 for mentions larger than 5). Examples contain assigned label as subscript and correct label as superscript.

Men. size	#SL errors	#RL errors	Example
2	36	28	[[Hong Kong] Disneyland]$^{\text{FACILITY}}_{\text{GPE}}$
3	15	5	[[Bank of China] Tower]$^{\text{FACILITY}}_{\text{ORG}}$
4	2	0	[[one hundred and twenty] dollar]$^{\text{MONEY}}_{\text{CARDINAL}}$

As it can be observed, both models make very few mistakes of this type (the data contains approximately 30K mentions). We are however only extracting cases involving sub-mentions seen in the data and might be underestimating the frequency of these errors. It may also be the case that given large annotated data of high quality, models make much fewer label bias errors. Investigating the several cases for which the RL corrects the supervised model, for the most cases, although not all, the sub-mention is correctly labeled, label which is further revised as the mention is processed. The word-level labels continue to be, as suggested by Table 4, mostly random, often rare labels such as LANGUAGE (LANG). Some examples include: [[Hong$_{\text{LANG}}$ Kong$_{\text{ORD}}$]$_{\text{GPE}}$ [International$_{\text{LANG}}$ Airport$_{\text{LANG}}$]$_{\text{FAC}}$]$_{\text{FAC}}$, [[[one hundred]$_{\text{CARD}}$ [and twenty]$_{\text{ORG}}$]$_{\text{CARD}}$ dollars]$_{\text{MONEY}}$, [Asia [- [Pacific [Special [Olympic Games]$_{\text{EVENT}}$]$_{\text{EVENT}}$]$_{\text{EVENT}}$]$_{\text{EVENT}}$]$_{\text{EVENT}}$ (omitting word labels)

5 Conclusions and future work

We have proposed a transition-based formulation for the mention detection task which generates mention parses as labeled binary trees and allows the model to create sub-mentions and flexibly revise labeling decisions. We have used reinforcement learning to train this model, by rewarding complete mention labels consistent with gold-standard annotations. In this work we have used a convolutional layer to build compositional representations for phrases. In the future we plan to investigate some of the methods proposed in the vast existing literature on this topic, including ways to pretrain such composition functions using unsupervised, language-learning-type of objectives.

We plan to also investigate ways to improve learning by bootstrapping with a policy learned in a supervised fashion, used to pretrain a critic in an actor-critic model.

Finally, we want to investigate different ways of assigning rewards. In particular, we plan to use rewards as a means to integrate training signal that might not be available (or may be too expensive to compute) at decode time, such as rewarding partial mentions which are consistent with syntactic constituents, grouping together words which exhibit high co-occurrence statistics, or mentions that agree with other mention detection systems.

References

Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. Globally normalized transition-based neural networks, 2016. URL `http://arxiv.org/abs/1603.06042`. cite arxiv:1603.06042.

A. Borthwick, J. Sterling, E. Agichtein, and R. Grishman. Exploiting diverse knowledge sources via maximum entropy in named entity recognition, 1998.

Léon Bottou. *Une Approche théorique de l'Apprentissage Connexionniste: Applications à la Reconnaissance de la Parole*. PhD thesis, Université de Paris XI, Orsay, France, 1991. URL `http://leon.bottou.org/papers/bottou-91a`.

Jason P. C. Chiu and Eric Nichols. Named entity recognition with bidirectional lstm-cnns. *CoRR*, abs/1511.08308, 2015. URL `http://arxiv.org/abs/1511.08308`.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *J. Mach. Learn. Res.*, 12:2493–2537, November 2011. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1953048.2078186.

Hal Daumé, III and Daniel Marcu. A large-scale exploration of effective global features for a joint entity detection and tracking model. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, HLT '05, pages 97–104, Stroudsburg, PA, USA, 2005. Association for Computational Linguistics. doi: 10.3115/1220575.1220588. URL http://dx.doi.org/10.3115/1220575.1220588.

R. Florian, H. Hassan, A. Ittycheriah, H. Jing, N. Kambhatla, X. Luo, N Nicolov, and S Roukos. A statistical model for multilingual entity detection and tracking. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 1–8, 2004.

Jiarong Jiang, Adam Teichert, Hal Daumé III, and Jason Eisner. Learned prioritization for trading off accuracy and speed. In *Advances in Neural Information Processing Systems 25*, pages 1331–1339, Lake Tahoe, NV, December 2012. URL http://cs.jhu.edu/~jason/papers/#jiang-et-al-2012-nips.

J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML2001*, 2001a.

John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, ICML '01, pages 282–289, San Francisco, CA, USA, 2001b. Morgan Kaufmann Publishers Inc. ISBN 1-55860-778-1. URL http://dl.acm.org/citation.cfm?id=645530.655813.

Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. Neural architectures for named entity recognition. *CoRR*, abs/1603.01360, 2016. URL http://arxiv.org/abs/1603.01360.

Francis Maes, Ludovic Denoyer, and Patrick Gallinari. Structured prediction with reinforcement learning. *Mach. Learn.*, 77(2-3):271–301, December 2009. ISSN 0885-6125. doi: 10.1007/s10994-009-5140-8. URL http://dx.doi.org/10.1007/s10994-009-5140-8.

NIST. The ACE evaluation plan. www.nist.gov/speech/tests/ace/index.htm, 2003.

Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Hwee Tou Ng, Anders Björkelund, Olga Uryupina, Yuchen Zhang, and Zhi Zhong. Towards robust linguistic analysis using ontonotes. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 143–152, Sofia, Bulgaria, August 2013. Association for Computational Linguistics. URL http://www.aclweb.org/anthology/W13-3516.

L. Ramshaw and M. Marcus. Exploring the statistical derivation of transformational rule sequences for part-of-speech tagging. In *Proceedings of the ACL Workshop on Combining Symbolic and Statistical Approaches to Language*, pages 128–135, 1994.

Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *International Conference on Learning Representations (ICLR)*, 2016.

Richard S. Sutton and Andrew G. Barto. *Introduction to Reinforcement Learning*. MIT Press, Cambridge, MA, USA, 1st edition, 1998. ISBN 0262193981.

Erik F. Tjong Kim Sang. Introduction to the conll-2002 shared task: Language-independent named entity recognition. In *Proceedings of CoNLL-2002*, pages 155–158. Taipei, Taiwan, 2002.

Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In Walter Daelemans and Miles Osborne, editors, *Proceedings of CoNLL-2003*, pages 142–147. Edmonton, Canada, 2003.